

This application is submitted in the name of inventors Jane Jin, Jie Chu, Maria Dos Santos, Shuxian Lou and Shujin Zhang, assignors to Cisco Technology, Inc., a California Corporation.

5

## SPECIFICATION

### TITLE OF INVENTION

### USE OF PRECEDENCE BITS FOR QUALITY OF SERVICE

### BACKGROUND OF THE INVENTION

10

#### Field of the Invention

The present invention relates to the field of data communications networks. More particularly, this invention relates to a method and apparatus for implementing a quality of service (QoS) policy in a data communications network so as to thereby prioritize network traffic into a plurality of service levels and provide preferential treatment of different classes of data traffic on the data communications network. A number of priority levels may be implemented in accordance with the invention.

15

#### The Background

20

This invention relates to switched packet data communications networks. There are a number of different packet types which are used in modern switched packet data communications networks.

25

FIG. 1A depicts a generic packet 8 using Layer 2 encapsulation. A number of

different Layer 2 encapsulation protocols are recognized. Each includes a MAC (media

access control) destination address 10 and a MAC source address 12. The data 14 may include Layer 3 encapsulated packet information. A CRC (cyclic redundancy check) 16 is also provided at the end of the Layer 2 encapsulation. The unlabeled block 18 may include an Ethernet type for Ethernet V 2.0 (ARPA) packets. The Ethernet type may 5 include IPv4 (IP) (in the future, IPv6), IPX, AppleTalk, DEC Net, Vines IP/Vines Echo, XNS, ARP or RARP. Other known encapsulations include SAP, SAP1, SNAP and the like. The meaning of the bits in block 18 differs among the different encapsulation protocols. This information is sometimes referred to as the Layer 2 Flow Information.

10 FIG. 1B depicts a typical IP packet (Layer 3) structure. The IP packet format is presently one of the most common Layer 3 packet types. The fields of importance to this disclosure are the "ToS value" or type of service 26 which is a preferably 8-bit field also known as Differentiated Services ("DS"), "prot-typ" or IP protocol type 28 (typically either TCP (transmission control protocol) or UDP (user datagram protocol), the Source 15 IP address 30 (usually the IP address of the originating station), the Destination IP address 32 (usually the IP address of the ultimate destination station), the Layer 4 source port number 34 (available for TCP and UDP packets only) and the Layer 4 destination port number 36 (available for TCP and UDP packets only). The Layer 4 flow information may be used to identify a particular packet flow as being the product of 20 (source port) or directed to (destination port) a particular application. The ToS /Differentiated Services field is used by routers of the data communications network to provide priority/delay/dropping services.

As the use of data communications networks increases worldwide, congestion of those networks has become a problem. A given data communications network, a given node on a data communications network, or a given link connecting two nodes has a certain capacity to pass data packets and that capacity cannot be exceeded. When data traffic on the data communications network becomes heavy enough that one can anticipate congestion problems which might lead to packets being dropped, it is desirable to implement a "Quality of Service" or QoS policy so as to give priority to certain types of traffic and/or customers so as not to drop their packets, thus assuring that critical communications are able to pass through the data communications network and/or providing a guaranteed service level, albeit at the expense of less critical communications.

The QoS level mapped into the ToS/Differentiated Services precedence bits (or other bits of the packet) is used by devices such as routers, gateways and switches within a data communications network in a number of ways. For example, packet rate limiting may be imposed to cap a user's throughput of packets to a given contract level specified in a service agreement with the user. In this way limited system services may not be over-used by those paying the least for those services. In another example, packet rate limiting may be imposed under congested network conditions to drop or modify packets based, at least in part, upon the QoS level associated with the packet. Those with a "higher" QoS level will do better in getting their communications through in congested conditions than those with a "lower" QoS level.

One way in which the ToS bits of the IP header are presently used is to use edge routers 38 as shown in FIG. 2 to interface between users 40 and core routers 42, 44, 46. In accordance with this approach the edge routers are used to set the ToS precedence bits (the highest order three bits of the eight-bit ToS/Differentiated Services field) based 5 upon the source IP address of the packet. While this can be effective in some environments, it is rather inflexible since it is based solely on pre-programmed IP addresses and creates administrative burdens in programming all of the edge routers. It is also ineffective in handling many types of roaming users and/or users with dynamically assigned IP addresses such as those obtained or leased from DHCP (dynamic host control 10 protocol) servers.

Accordingly, a need exists for a QoS system which can provide for the QoS needs of roaming or DHCP users while simplifying the task of programming QoS levels associated with the user. Furthermore, it would be desirable to provide a QoS system that 15 provides on-demand QoS rather than a fixed QoS for each user to better accommodate the shifting needs of users.

#### SUMMARY OF THE INVENTION

Certain bits of a packet, such as bits in the IP header of an IP packet, are used to 20 designate the type of service or Quality of Service (QoS) level to be afforded to the packet as it passes through a data communications network. A user entitled to a certain QoS level logs into a service selection gateway SSG. The SSG queries an authentication, authorization and accounting (AAA) server in response to a log-in attempt by the user.

Upon authorization, the AAA server returns an access accept signal in addition to an indication from the user's service profile (user profile) as to the QoS level to be afforded the user. While the user is logged in, all packets are routed through the SSG. The SSG sets the certain bits of the packet in accordance with the user's assigned QoS level so that 5 as the packets are routed through the data communications network, they are consistently afforded the assigned Quality of Service level. In another aspect of the invention, on-demand QoS may be provided by the SSG and accounted for by communications with the AAA server.

10

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a diagram showing the structure of a typical Ethernet packet.

FIG. 1B is a diagram showing the basic structure of a Layer 3 IP packet.

15

FIG. 2 is a system block diagram of a typical network interface to a user.

FIG. 3 is a system block diagram of a data communications network in accordance with a presently preferred embodiment of the present invention.

20

FIG. 4 is a flow diagram of a first method of setting QoS for an IP packet in accordance with one embodiment of the present invention.

FIG. 5 is a flow diagram of a second method of setting QoS for an IP packet in accordance with another embodiment of the present invention.

**DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT**

5 Those of ordinary skill in the art will realize that the following description of the present invention is illustrative only and not in any way limiting. Other embodiments of the invention will readily suggest themselves to such skilled persons after a perusal of the within disclosure.

10 In accordance with a presently preferred embodiment of the present invention, the components, processes and/or data structures are implemented using a gateway device implemented using C++ programs running on an Enterprise 2000™ server running Sun Solaris™ as its operating system. The Enterprise 2000™ server and Sun Solaris™ operating system are products available from Sun Microsystems, Inc. of Mountain View, California. Different implementations may be used and may include other types of operating systems, computing platforms, computer programs, firmware and/or general purpose machines. In addition, those of ordinary skill in the art will readily recognize that devices of a less general purpose nature, such as hardwired devices, devices relying on FPGA (field programmable gate array) or ASIC (Application Specific Integrated Circuit) technology, or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein.

15

20

In accordance with a presently preferred embodiment of the present invention, the three Precedence bits (the three highest order or most significant bits of the 8-bit Type of Service (ToS)/Differentiated Services field 26 of the IP packet header) are used to designate the type of service or Quality of Service (QoS) level to be given to packets sourced from a particular user. Those of ordinary skill in the art will realize that the particular bits used are not particularly critical, for example, the CoS (Class of Services) bits of an IEEE 802.1q encapsulated frame could be used as could the CoS bits in an ISL frame. Other bits or fields could also be designated to carry the QoS level information. A 3-bit ToS permits up to 8 levels of service. Initially, more than three bits could be used, if desired.

The present invention makes use of a user's service profile. A service profile, sometimes referred to as a user profile, contains information relating to a particular user's network access account. For example, it may include an identification of the user's last known home PoP (point of presence) or home gateway located in a PoP. It may include the identification of one or more domain name server(s) (DNS) to use in resolving domain names to IP addresses. It may include details about the user's service agreement with the ISP (internet service provider) servicing the user's account. Such information may include an identification of the QoS level to be provided to the user. For example, if the user is a "Platinum" user, he or she might receive a level 8 ToS thus granting his/her packets the best service. A normal user not paying a premium for a higher QoS might be afforded a level 1 ToS.

The user profile is stored in AAA (authentication, authorization and accounting) servers disposed in various locations in the data communications network.

The authentication, authorization and accounting (AAA) server performs user authentication, user authorization and user accounting functions. It may be a Cisco ACS or product such as Cisco Secure, available from Cisco Systems, Inc. of San Jose, California, or an equivalent product. In accordance with a presently preferred embodiment of the present invention, the Remote Authentication Dial-In User Service (RADIUS) protocol is used as the communication protocol between the service selection gateway and the AAA server. RADIUS is an Internet standard track protocol for carrying authentication, authorization, accounting and configuration information between a gateway device that desires to authenticate its links and a shared AAA server. Those of ordinary skill in the art will realize that other Internet protocols such as TACACS+ can be used as acceptable authentication communications links between the various communications devices that encompass the data communications network and still be within the inventive concepts disclosed herein.

The service selection gateway (SSG) is a device which couples the user to the data communications network. The user may log directly into the service selection gateway or may do so via an intermediate server such as a network access server. The term gateway is not meant to be limited to a single type of device, as any device, hardware or software, that may act as a bridge between the user and the network may be considered a gateway for the purposes of this application. In accordance with a presently

preferred embodiment of the present invention, the service selection gateway is a model 6510 service selection gateway available from Cisco Systems, Inc. of San Jose, California.

5       Turning now to FIG. 3, when the user 48 logs-in to the SSG 50, either directly or through one or more intermediate devices, the SSG 50 obtains the user's service profile from an AAA server 52. The user profile will contain an additional field detailing the QoS level to be afforded the user in accordance with a presently preferred embodiment of the present invention. The SSG 50 stores the QoS level associated with the user in a local  
10      memory or cache associated with the user and, in effect, acts as an edge router for the user in all subsequent outbound communications forwarded to the Internet 54 or other data communications network during the session. The SSG 50 simply replaces the ToS/Differentiated Services with a bit pattern corresponding to the user's QoS level.

15      This approach provides a number of important advantages. First, the QoS level for the user need only be set once in establishing the user's service profile. The existing AAA system will assure that the user's service profile is available regardless of the PoP that the user logs-in on. This means that roaming users will be supported where they could not easily be supported in the edge router system approach. Second, maintaining  
20      and modifying the QoS level associated with a user is quite simple--simply modify the user service profile once. Third, this approach allows implementation of a pay-per-use approach to QoS where a user may negotiate with an SSG to obtain a particular QoS level that the user would not otherwise receive to cover a particular session or one or more

particular flows within a session. This could be used where a user desires to send time critical communications such as voice over IP or video over IP where dropouts are distracting and undesirable. In such a case, the user could be billed or in some way charged for the use, if desired. The SSG in conjunction with the AAA server can 5 effectively track the user's use of the on-demand service and send RADIUS (or equivalent) packets to the AAA server. A number of billing and use schemes can now be implemented that were unavailable under the edge router system.

Turning now to FIG. 4, a method in accordance with one embodiment of the 10 present invention is described. At reference numeral 100 the QoS level to be afforded all packets sent by the user is set in the user profile of the user. This is stored in AAA servers preferably distributed about the data communications network.

At reference numeral 102 the user attempts log-in at a PoP containing a service 15 selection gateway--either directly or through a network access server or other intermediate server.

At reference numeral 104 the SSG queries an AAA server using the RADIUS (or an equivalent) protocol and obtains the QoS level specified in the user's service profile.

20

At reference numeral 106 the SSG applies the QoS level for the user by setting the ToS precedence bits of each IP packet sent by the user to the QoS value specified in the user's service profile or to a direct mapping of that value. Note that QoS level may be

implemented using other bits in the packet as would be apparent to those of ordinary skill in the art.

5 The above procedure describes a "subscription"-based QoS approach. A different "pay-per-use" QoS approach is detailed in FIG. 5.

At reference numeral 108 the user attempts log-in at a PoP containing a service selection gateway--either directly or through an intermediate network access server or other intermediate device.

10

At reference numeral 110 the user signals the SSG of a selected QoS level to use. This may be for all packets sent during the session, for all packets sent during the session of a particular protocol type or from or to a particular Layer 4 port number (e.g., voice over IP might be set to a higher QoS than standard traffic) and the like. This can be 15 accomplished through a conventional menuing system, through a dashboard application program in communication with the SSG and/or through direct command to the SSG.

15

At reference numeral 112 and SSG optionally sends accounting packets detailing the user's use of the QoS level within the data communications network so that pay-per- 20 use may be realized.

20

At reference numeral 114 the SSG applies the user selected QoS levels to selected packets in accordance with policy set by the user to the ToS precedence bits of the ToS

field of IP packets sent by the user. This happens until the user selects another QoS level or logs-out.

#### Alternative Embodiments

5 While embodiments and applications of the invention have been shown and described, it would be apparent to those of ordinary skill in the art, after a perusal of the within disclosure, that many more modifications than mentioned above are possible without departing from the inventive concepts herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

100 99 98 97 96 95 94 93 92 91 90 89 88 87 86 85 84 83 82 81 80 79 78 77 76 75 74 73 72 71 70 69 68 67 66 65 64 63 62 61 60 59 58 57 56 55 54 53 52 51 50 49 48 47 46 45 44 43 42 41 40 39 38 37 36 35 34 33 32 31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1